Three Mathematical Views of *In C*

Donald Spector
Department of Physics
Hobart & William Smith Colleges
Geneva, NY 14456 USA
spector@hws.edu

Abstract

This year marks the fiftieth anniversary of Terry Riley's seminal composition *In C*, a groundbreaking piece that led to the minimalism movement in music. Strikingly, *In C* also presents a tangible embodiment of several abstract mathematical structures. This essay explores three mathematical perspectives on *In C*. The first is that of diffeomorphisms and differential geometry; the second is functional integration; and the third is a strange hybrid of finite state automata, Markov chains, and Conway's Game of Life. Through Riley's composition, listeners can directly experience abstract mathematical ideas in these three arenas. In addition, the piece suggests new ways to look at finite state automata.

Introduction

The 1960s saw the emergence of many new strands of music: electronic music, computer music, various rock music genres, and minimalism. Minimalism developed under the radar; "[t]he term itself, now common currency, appeared in the mid-sixties but was largely unheard outside of art and avant-garde music circles until the eighties" [**?**]. Yet, while electronic and computer music seemed to capture the zeitgeist of the era, and the evolution of rock music drew popular attention, it was minimalism that emerged as the dominant genre of classical experimental music over the decades that followed. Compositions such as Steve Reich's *Four Organs* (1970) and *Music for 18 Musicians* (1976), Tom Johnson's *An Hour for Piano* (1971), the Philip Glass/Robert Wilson opera *Einstein on the Beach* (1976), and the ambient music of Brian Eno and Harold Budd (1978-1984) established minimalism as an ongoing presence in the world of music [**?**] [**?**], leading to such influential contemporary minimalist composers as John Adams and Louis Andriessen.

The piece that is generally credited with launching the minimalist movement in music is Terry Riley's *In C* (1964). It was not the first minimalist piece—there are certainly earlier examples outside the Western tradition [**?**], and even in Europe and the US, there were precursors such as Lamonte Young's music [**?**]—but *In C* is recognized as the first minimalist piece that struck a nerve and got widely noticed [**?**]. Thus, while Philip Glass, Steve Reich, and Brian Eno are better known, the lineage of their works traces back to Terry Riley and *In C*.

It turns out that *In C* is not only significant musically, but also mathematically. As we will see below, the structures of *In C* contain embodiments of mathematical ideas that are seldom seen outside the halls of research. As such, the piece broke new ground in more than aesthetic and stylistic ways.

In C: A Description of the Piece

In C consists of a set of 53 melodic patterns of music, along with directions on how the musicians are to move through those 53 musical fragments. Here are the highlights of the performing directions, taken from the score [**?**].

- "All performers play from the same page of 53 melodic patterns played in sequence."
- "The tempo is left to the discretion of the performers, obviously not too slow, but not faster than the performers can comfortably play."
- "The ensemble can be aided by the means of an eighth note pulse played on the high c's of the piano or on a mallet instrument. It is also possible to use improvised percussion in strict rhythm (drum set, cymbals, bells, etc.), if it is carefully done and doesn't overpower the ensemble. All performers must play strictly in rhythm and it is essential that everyone play each pattern carefully. It is advised to rehearse patterns in unison before attempting to play the piece, to determine that everyone is playing correctly."
- "Any number of any kind of instruments can play. A group of about 35 is desired if possible but smaller or larger groups will work."
- "Patterns are to be played consecutively with each performer having the freedom to determine how many times he or she will repeat each pattern before moving on to the next."
- "There is no fixed rule as to the number of repetitions a pattern may have, however, since performances normally average between 45 minutes and an hour and a half, it can be assumed that one would repeat each pattern from somewhere between 45 seconds and a minute and a half or longer."
- "It is important not to hurry from pattern to pattern but to stay on a pattern long enough to interlock with other patterns being played. As the performance progresses, performers should stay within 2 or 3 patterns of each other. It is important not to race too far ahead or to lag too far behind."
- "IN C is ended in this way: when each performer arrives at figure #53, he or she stays on it until the entire ensemble has arrived there. The group then makes a large crescendo and diminuendo a few times and each player drops out as he or she wishes."

What we have here is a piece in which a fairly large collection of performers moves through the same 53 musical fragments, doing so at a common tempo that is rigorously adhered to through the piece, but with each performer spending as long or as short a time on each fragment as he/she desires by repeating it as many times as he/she desires. The performers all start and end at the same point and pass through the same 53 way stations, but their journeys are distinct. Thus we have a framework within which a multitude of related individual performances occur simultaneously. The relationships among these varied individual performances embody several mathematical ideas, which we explore below.

In C and Diffeomorphisms

One of the central themes of geometry is how we can use measurements of distance on a space to characterize that space. One particularly useful way to probe this is to consider transformations of a space, and to study how those transformations affect distances. This approach also provides insights into the structure of *In C*. I will first develop the notion of transformations in geometry and identify their analogues in music, and then consider diffeomorphisms and *In C*.

The geometry with which most people are familiar, Euclidean geometry, is the geometry of flat (i.e., uncurved) space. Among the operations with which we can transform flat space are the so-called rigid transformations, that is, transformations that leave the distances between points unchanged [**?**]. For example, if we take a sheet of paper and rotate it or take a sheet of paper and move it left or right, the distances between points do not change. These two types of transformations—rotations and translations—are not, however, the only kinds of transformations of interest.

Scale transformations, which stretch or shrink distances uniformly over the whole space, may also be used to transform a space. Transformations of this type are executed when a photocopier takes an original and makes a re-scaled image as the copy, or when we zoom in on an image on a computer screen.

Translations—the shifting of a space by some fixed displacement—appear musically in rounds or perpetual canons. Two or more musicians play the same notes, but start doing so at different times. Thus the different parts are shifted from each other in time, but the relative distance in time between any two notes is the same in each part. Of course, not any piece is aesthetically amenable to being turned into such a canon. This is a second, implicit criterion: making a musical embodiment of a mathematical structure is only of interest when the result is musically satisfying.

We also see the musical equivalent of a shift when a piece of music is transposed from one key to another. Visually, this shift is up or down the page in the score. Tonally, the shift is up or down in tone. However, if we want to think of this as a simple shift transformation, we need to define a proper notion of distance. Initially, one might imagine that this refers to frequency, but in fact differences in frequency are not preserved under transposition; ratios of frequencies are. In order to have something that behaves like a distance, we can refer to notes by the logarithms of their frequencies. Thus, for example, if we transpose a piece by an octave, every note doubles in frequency, and thus the difference between the logarithms of the frequencies of every original note and its corresponding transposition is the logarithm of 2.

A fugue generally begins with a combination of both kinds of translation, with each voice coming in later than the original voice and transposed, e.g., up a fifth, although as the fugue develops, the different musical lines are not fixed to be identical to each other (indeed, small changes can even appear in the first reiteration of the initial theme).

What about scale transformations? The simplest embodiment of these in music is through the choice of tempo. Geometrically, scale transformations preserve ratios of distances, even though individual distances will not be preserved. If we consider the notes of a piece spread out over time, then adjusting the tempo has the same net effect: the lengths of individual notes will change, but the ratios between their lengths will stay the same (e.g., a half note will always last twice as long as a quarter note). This is used in augmentation and diminution canons, for example, in which the initial line occurs with note lengths cut in half or doubled, respectively. Tonally, transposition from one key to another also serves as a scale transformation, as long as we think of notes in terms of their frequencies (and not logarithms thereof, as we did above).

The possibilities within geometry became greater when we bring in the tools of analysis, leading to the field of differential geometry [**?**]. Fortunately, we will not need the full structure of differential geometry, but simply a consideration of a new type of transformation called a *diffeomorphism*.

Diffeomorphisms take the above sorts of transformations—rotations, re-scaling, and translations—but instead of applying them uniformly over the whole space, we allow the size of the associated transformations to differ by location. (There is also a continuity condition, but we will leave that technical point aside.) Thus, for example, imagine taking a sheet of rubber, marking a grid of points on it, and then stretching and distorting the piece of rubber in ways that vary across the sheet. The end result will be that the grid of points will no longer be arrayed at regular distances or orientations from each other. To see this, we could, after distorting the sheet of rubber, place a second grid of points on the sheet, and then measure the distance and orientation of the old grid points by using the metric structure of the new grid points; this would enable us to characterize the particular diffeomorphism that has been applied.

How would you represent this musically? Let us think about a piece of music unfolding in time, so it is a space labeled by a single dimension. The corresponding diffeomorphism transformations, then, are what you might get by taking a piece of a rubber band, and stretching it by different amounts at different points.

A simple way to handle this musically would be simply to change the tempo in various ways as you went through the line of music. The problem here is that with only a single line of music, you could not tell what had happened. Yes, you could compare different performances, but a single performance would not reveal what had happened; we could not tell if we were hearing a shifting tempo, or simply notes of varying durations.

If you tried to address this by taking two or more musicians playing the same set of notes, but each using tempi that differed over time in personal and distinct ways, the different lines would be diffeomorphically related to each other, but in an uncontrolled way that would be difficult for the average listener to discern, and challenging to set up to ensure an aesthetically pleasing (within the context of our cultural musical framework) experience.

What Terry Riley achieves in *In C* is to find a way to represent diffeomorphisms that is apparent to the listener upon a single hearing, is aesthetically pleasing, and is, most importantly from our point of view, embedded in the very structure of the piece.

We can imagine that *In C* is like 53 strips of elastic sewn together into one long strand. The first strip of elastic corresponds to musical fragment #1, the second to #2, and so on. Thus, the music played by any single performer is just the sequence of these elastic segments, one played after the other. So far, so good: every player has the same score to play.

Now is where the diffeomorphisms come in. Each performer can stretch or shrink the length of each of his/her segments separately—separately from how his/her other segments are stretched or shrunk and separately from how other performers are stretching or shrinking their segments. The duration is modified not by changing the beat, but simply by playing the fragment more or fewer times. In short, we view a single player's line of music as a geometrical line, approximated by having been divided into 53 segments. The number of times a musician plays a segment determines its geometrical length.

Thus, each player's performance is a diffeomorphically transformed version of one underlying musical line. (Strictly speaking, the discretization of the piece into 53 segments means that this is a discretized version of a diffeomorphism.) By having a set of musical fragments that each player moves through, but with all the musicians playing at a fixed tempo, Riley has found a solution to the problem of how to embody the analogue of diffeomorphisms in music while retaining a sound that is aesthetically pleasing. In addition, the regular rhythmic beat that accompanies the other musicians serves the function of the second array of grid points in our example above: they provide a framework with which we can compare the various musicians' performance choices, and auditorily experience the stretching and shrinking in time that each player adopts at different points in the piece, even upon a single hearing.

As a final note, it is worth commenting that general relativity, Eintein's theory of gravity, has as its fundamental principle that the laws of nature are unchanged under diffeomorphisms of spacetime [**?**][**?**]. Riley's piece embodies something of this invariance principle as well. Each player is given the same score, the same set of instructions. Thus each player's performance—while phenomenologically distinct—is a performance of the same musical score. Riley allows us to identify what is invariant amidst these diffeomorphism transformations.

In C and Path Integrals

Calculus is familiar far and wide. Given a function, derivatives allow us to study how that function varies, which leads, for example, to the standard rules for identifying extrema. Integration, on the other hand, allows us to sum together the behavior of a function over a collection of points; for example, the integral of a function divided by the length/area/volume of the region over which the integration is performed allows us to find the average value of a function.

But what happens when instead of having functions of real numbers, the inputs to our functions are entire paths? For example, we might ask the following question: Given that we are to design a slide that starts at one point and finishes at another, how should it be designed to minimize the length of a time it will take someone to slide from start to finish? This is clearly a minimization problem—what shape slide minimizes the sliding time—but the inputs are not numbers, but rather the various curves that go from the starting point to the ending point. The mathematical arena that allows the solution of such a problem is the calculus of variations, a topic that regularly and easily shows up in an undergraduate education [**?**].

But what about the equivalent generalization of integral calculus? This somehow is less well-known, but the relevant technique is called functional integration or path integration. In fact, there is a way, developed by Feynman, to formulate to quantum mechanics entirely in terms of functional integration [**?**]. This method is essential to modern studies of quantum field theory [**?**][**?**], and has led to some of the very few exact results in particle physics [**?**][**?**][**?**].

In a functional integral, one considers all paths satisfying some condition (e.g., a starting and finishing point within the space in question); this forms the region of integration. One then identifies a function of these paths, and sums together the value of the function over all these paths. The application of functional integrals as a formal tool, especially in particle physics and statistical mechanics, has extended even beyond the contexts in which the rigorous mathematical treatment has been worked out. Thinking of a physical process as the net result of a kind of average over all the different ways that process could unfold turns out to be a rich and powerful notion.

In the course of a mathematical education, one can formulate these ideas, but how does one convey such an idea without the benefit of coursework in functional analysis? One answer is to attend a performance of *In C*.

Let us consider how the music is generated in *In C*. A large-ish number of performers (ideally, around three dozen) all perform simultaneously. They are all taking paths through the same 53 melodic ideas, but each is taking his/her own path. Here, the set of allowed paths is not identified by where they are in space, but by how many times the performer plays each melodic form when he/she encounters it. Here, then, we are not focusing on the common underpinnings of these paths (as we did in the previous section), but instead understanding each musician's performance as a path that is a different instance from the set of allowed paths.

We have the paths; where is the integral? Consider now what the audience hears. Each performer has a path (e.g., one performer's path is to play fragment #1 seven times, fragment #2 three times, and so on), and this leads then to the sound that that performer produces. The sound the musician produces is a function of his/her path. But the audience does not just hear a single musician; the audience hears the collected sounds of all the musicians, and each musician's performance corresponds to a different path through the 53 fragments, a different selection from the space of all possible paths. Thus the actual outcome of the piece is a sum over all the performed paths through the 53 fragments, and so the piece has the musicians produce the experience of a sonic path integral for the audience.

We note that in specifying that a typical performance would have around three dozen performers, Riley is seeking to ensure that we are not experiencing a few individual musical lines, but rather are experiencing a truly collective sound, one that can only fully be appreciated in its collective form. There are enough performers that the average of all the performances is not especially affected by one performer or another; that is to say, we are in the regime where it is reasonable to think of the sum as an approximation of an integral. While he probably did not have the mathematical language to describe it this way, Riley was clearly trying to create for the audience the experience of a functional integral. This gives us, then, a second mathematical way to think about *In C*.

In C, Finite State Automata, Markov Processes, and the Game of Life

Our final approach to *In C* is to consider each performer playing the piece as executing an instance of a Finite State Automaton (FSA) [**?**], but in an unusual way that links it to an ensemble of other FSAs, following a protocol that bears certain similarities to Conway's Game of Life [**?**].

A finite state automaton [**?**] is a machine that can be in one of a finite number of states, and that, at each step in time, can change to another of these states. Let us label the state $A_1, A_2, \ldots, A_n$. If the machine is in a state A_k, then there is a list of possible subsequent states into which it can evolve at the next step, generally smaller than the full set of possible states. For example, it might be that after A_1 the system can either stay in A_1 or move to A_2 or A_3; in A_2, perhaps the system could next be in A_1, A_4, or A_7. Once we know all the possibilities, we can arrange the states in a graphical structure, with arrows showing which states can evolve into which other states.

One category of FSA consists of the deterministic FSA. In these, there are some conditions that indicate not simply which states can evolve into which, but in any given circumstance, which subsequent state will arise. The rule might refer to the environment or to some internal state variable.

A second category of FSA are the non-deterministic FSA. Markov chains provide one model of this [**?**]. Suppose, in our example above, we have that if the system is in state A_1, there is probability 0.4 that after one step in time it will stay in A_1, probability 0.5 that it will evolve to A_2, and probability 0.1 that it will evolve into A_3. Then at any time the system is in state A_1, a metaphorical roll of the dice will happen; the result of that roll and the associated probabilistic rules will determine what next state will appear. A non-deterministic FSA thus does not have a unique evolution, although we can discuss its expected evolution, and perform statistical analyses on how it will behave. (It is also the case that in an FSA the probabilities do not have to be fixed over time, as they are in a Markov chain.)

If you look at *In C*, we see that we can view each musician as executing an FSA. After playing melodic fragment #1, there are two choices: repeat melodic fragment #1 or move to melodic fragment #2. In general, after playing melodic fragment #*n*, there are two choices: repeat melodic fragment #*n* or move to melodic fragment #$(n+1)$. A diagram of the piece, then, as an FSA would be be a collection of 54 states, with each of the first 53 states—corresponding to the 53 musical fragments—having an arrow to itself or an arrow to the subsequent state. The final, 54th state is a "stop playing" state, at which point the musician and the FSA halt. (Since all players begin with melodic fragment #1, it is superfluous to add a "begin playing" state.)

Now is this a deterministic or a non-deterministic FSA? Presumably, there are some internal parameters of the musician's brain through which he/she decides whether to repeat or move on, so with some deep understanding of the brain, this might be a deterministic FSA. In practical terms, however, it is probably something closer to a non-deterministic FSA, since neither the performer nor the audience could in practice indicate ahead of time when a given performer will move on.

And yet, this is not a simple Markov process. One constraint is that Riley gives guidance as to typically how long a performance of *In C* should be, and thus approximately how long one should typically spend on one of the musical fragments. Consequently we have a situation in which the guidance suggests, indirectly, a set of probabilities for how many times the "repeat" loop will be executed for a given fragment. In fact, it would be possible to achieve an expected value for the duration of the performance by specifying a suitable set of fixed probabilities at each state, and so it is certainly possible that each player could be executing a straightforward Markov process, but it is unlikely that one would do so. After a certain point, with each repetition of fragment #*n*, it becomes likelier that the performer will move on to fragment #$(n+1)$.

But that is not all. Riley gives specific instructions that cause the FSAs associated with the various musicians to influence each other, and thus influence the corresponding probabilities we might use to describe the behavior of a single musician's FSA. Musicians are told to stay within two or three musical fragments of their associates. Thus, for example, once one musician/FSA moves to fragment #5, this triggers a reaction in which the likelihood that other musicians/FSAs still on fragment #1 will move on to #2 increases. Likewise, if there are still many musicians/FSAs at or around fragment #2, the likelihood that the musician at fragment #5 will move on to #6 is reduced. In short, the behavior of the individual is tied to the behavior of the aggregate, and the behavior of the aggregate is tied to the behavior of the individuals. In a broad sense (but certainly not in the particulars), this is somewhat like Conway's Game of Life [**?**], in which the fate of a given square depends on what the fate of its neighbors has been, and the fate of the neighbors depends on part on what the fate of that given square has been. (One also sees similar behavior naturally when birds flock or insects swarm.) In *In C*, we do not have a notion of neighbors in the same way (although a musician is probably more likely to hear whether a physically adjacent musician has moved on or not than than whether one on the other side of the room has or has not done so), but we do have an interesting coupling of these FSAs. Strikingly, by and large, the mathematical behavior of FSAs that are coupled in this way has not received particular attention; having now encountered this idea in Terry Riley's music, however, I have begun to think about how such ensembles of FSAs might be productively formulated and analyzed. So just as Terry Riley helped pave new ground in music in 1964, a mathematical consideration of his work leads to new arenas deserving of mathematical investigation half a century later. It is hard to imagine a richer outcome.

Conclusion

It is certainly a familiar situation in the arts for a work of art to have more than one artistic interpretation. What is striking here is that we see that same richness in terms of a work of art's mathematical interpretation. *In C*, which changed the face of music, did not just represent the embodiment of new musical ideas; it was a piece whose very structure specifically encodes in a tangible way abstract ideas from the world of mathematics. Ideals like diffeomorphisms and path integrals can be hard to grasp from a theoretical perspective without serious training in mathematics. But *In C* makes these mathematical notions experientially accessible. As a new kind of musical composition, it embodies mathematical constructs that had not previously been embodied directly in a single piece of music. Indeed, this newness reflects back the other way: once we understand *In C* as asking the musicians to become FSAs, we see a new paradigm of an ensemble of coupled FSAs emerging, FSAs whose behavior is probabilistic but depending on the behavior of their associated FSAs. *In C* invokes a structure in the space of FSAs similar to, but more intricate than, that in Conway's Game of Life, and in doing so, opens up new questions to examine regarding FSAs. This is the true richness of the intersection of mathematics and the arts: the intersection not only gives us a new language to talk about each field, it opens up new doors of investigation in each field. The music might be minimal in style, but its impact is anything but.

References

[1] Mary L. Boas. *Mathematical Methods in the Physical Sciences, 3rd edition*. Wiley 2005.

[2] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory, 2nd edition*. Wiley Interscience 2006.

[3] Richard P. Feynman and A.R. Hibbs. *Quantum Mechanics and Path Integrals*. McGraw-Hill 1965. Emended addition Dover 2005, with material by Daniel Styer.

[4] Martin Gardner. "Mathematical Games: The fantastic combinations of John Conway's new solitaire game 'life'." Scientific American 223 (1970) 120.

[5] Mark V. Lawson. *Finite Automata*. CRC Press LLC 2004.

[6] Wim Mertens. *American Minimal Music.* Kahn & Averill 1983.

[7] Charles W. Misner, Kip S. Thorne and John Archibald Wheeler. *Gravitation*. W.H. Freeman & Co. 1973.

[8] S. Nussinov, "Baryon-meson mass inequality," Phys. Rev. Lett. 51 (1983) 1081.

[9] Barrett O'Neill. *Elementary Differential Geometry, revised 2nd edition*. Academic Press 2006.

[10] Michael E. Peskin and Dan V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press 1995.

[11] Keith Potter. *Four Musical Minimalists: La Monte Young, Terry Riley, Steve Reich, Philip Glass*. Cambridge University Press 2000.

[12] Pierre Ramond. *Field Theory: A Modern Primer*. Benjamin/Cummings 1981(2nd edition, 1989).

[13] Terry Riley. *In C*. G. Schirmer/Associated Music Publishers 1964.

[14] Kiran Sande. *Less Is More: A Brief Survey of Minimalism*. http://www.redbullmusicacademy.com/magazine/less-is-more-a-brief-survey-of-minimalism (accessed 1 February 2014).

[15] Edward Strickland. *Minimalism: Origins*. Indiana University Press 1993.

[16] Steven Weinberg. *Gravitation and Cosmology*. Wiley 1972.

[17] Don Weingarten, "Mass Inequalities for QCD," Phys. Rev. Lett. 51 (1983) 1830.

[18] E. Witten, "Some Inequalities among Hadron Masses," Phys. Rev. Lett. 51, (1983) 2351.